University of Illinois at Chicago
Fall 2018

CS 521 — Statistical Natural Language Processing (Current Topics in NLP)

Course Syllabus

Room: BH 317
Time: Tu Th 9:30-10:45
Class Web Site: via Blackboard; discussion on Piazza

Instructor: Barbara Di Eugenio
Office: 916 SEO (in transition to 1138 SEO)
Phone: 6-7566 (may change)
E-mail: bdieugen@uic.edu
Office Hours: Tue & Th, 2-3pm (starting week of Sep 4)

Course Objectives

The aim of this course is to cover some of the empirical methods that are widespread in Natural Language Processing (NLP), and to provide an overview of current, cutting-edge research in the field. This course is meant to provide students with both the foundations necessary to understand cutting edge research papers that use these techniques in any area of NLP; and with an in-depth and critical look at some of the "hottest" applications that use those techniques, for example semantic inferencing, question answering, opinion mining and sentiment analysis, dialogue management, summarization, language in social media, human-robot interaction.

At this point, a more appropriate name for this course would be *Current Topics in NLP*, since 99% of NLP these days uses empirical methods.

Reading Materials

1. **Required Textbook / Readings:** Daniel Jurafsky and James H. Martin. Speech and Language Processing, second edition / draft third edition. Prentice Hall, 2008 / 2019. SNLP 3rd edition

2. **Additional materials** from: Christopher Manning and Hinrich Schütze. Foundations of Statistical Natural Language Processing. The MIT Press, 1999

3. **Many Articles** from the literature. ACL Anthology

Prerequisite

CS 421 (Introduction to NLP) or equivalent, or permission of the instructor. If no background on NLP, you may be allowed in the class **only if** you have foundations in at least one of Artificial Intelligence, Information Retrieval, Machine Learning / Data Mining. Please discuss your background with the instructor.

IMPORTANT. Students allowed in the class who **don't have** NLP background, or whose NLP background is rusty, **must study** the following in the textbook on their own (online): sec. 3.1-3.3, 3.4.1 (n-grams); 8.1-8.4 (POS tagging, HMM); 10.1-10.4; 11.1-11.4; 12.1-12.3 (ch. 10-12 are on grammars and syntactic parsing); ch. 18 (semantic roles); 25.1-25.2 (dialogue systems)
NOTE: because the third edition on-line is work in progress, chapter numbers may change. Please check with instructor if topic of chapters doesn't align with short descriptions above.

Important Note: Laptop Usage in class

I don't mind if you use your laptop / tablet in class but its usage must be related to class – ie taking notes. Hearing constant typing is distracting to the instructor and classmates. I reserve the right to ask you to close it down if I find it disruptive.

Tentative Schedule

Dates	Topic	Readings
	Weeks 1-6: Foundations	
8/28, 8/30, 9/4	Introduction, n-grams, smoothing	Ch. 3 (new edition)
9/6	Tagsets, Intercoder Agreement, Crowdsourcing	
9/11, 9/13	HMMs	Ch. 5, 6 (old edition)
9/18, 9/20	Information Theory, Max. Entropy, MEMMs	
9/25	Prof Petrov	IBM Watson
9/27, 10/2, 10/4	Word2Vec & Neural Networks	Ch. 6-7 (new edition)
10/9 (Week 7)	Catch-up, Review	
10/11	**Midterm Exam**	
	Weeks 8-14: Applications (Tentative!)	
Week 8-9	Language in Social Media	
Week 10-11	Dialogue Processing, Multimodal Interaction	
Week 12-13	Semantic Inferencing, Question Answering	
Week 14	BioNLP, NLP for health sciences	
Week 15	Catch up, and / or Project Presentations	

From week 7 on, we will read articles from the literature. Each student will be asked to present 1 or 2 papers, and to be the discussant for 2 or 3 other papers – this means writing a short written critique for the paper in question and be ready to participate in discussion. Exact workload will depend on class size.

Important Dates

Date	Event
Th 10/4	Project Proposal (One page)
Th 10/11	Midterm (Open Book and/or Notes)
Th 11/8	Project "in progress" report (One page)
Week of 12/10 (Finals Week)	Project Presentations and Submission

Grading Criteria

The class will be graded as follows:

- **Midterm**: 25%.
- **Paper Presentation(s)**: 10% each.
- **2-3 Paper Critiques**: 5% each.
- **Class Participation**: 2%
- **Project**: The remaining points, divided as follows:
  - project originality, work execution, thoroughness, etc: half of project points;
  - proposal + intermediate report $\frac{1}{8}$; final presentation $\frac{1}{8}$; final report $\frac{1}{4}$.

  The project should be done in pairs.

Possible readings

Here's a list of potential papers we may discuss, or conferences where papers on the topic are presented. It's neither an exhaustive list, nor does it mean we will read all (or any!) of these papers: it provides you with some ideas about the kind of papers we will be looking at.

Social Media Language Processing [Agarwal *et al.*, 2013; Di Eugenio *et al.*, 2013; Prabhakaran and Rambow, 2013; Bracewell, 2015; Pavlick and Tetreault, 2016; Goldwasser and Zhang, 2016; Misra *et al.*, 2017; Preoţiuc-Pietro *et al.*, 2017; Ghosh *et al.*, 2017; Johnson and Goldwasser, 2018]; papers from workshops on Language in Social media e.g.
`http://aclweb.org/anthology/W/W17/#1100`,
`https://aclanthology.coli.uni-saarland.de/events/ws-2018#W18-35`,
First workshop on trolling aggression and cyberbullying

Conversational Agents and Dialogue Processing [Singh *et al.*, 2002; Mairesse and Walker, 2010; Chen and Di Eugenio, 2013; Janarthanam *et al.*, 2013; Fang *et al.*, 2014; Manuvinakurike *et al.*, 2016; Shalyminov *et al.*, 2017; Zhang *et al.*, 2018]
Papers from SIGDIAL (ACL-ISCA Special Interest group on Discourse and Dialogue) e.g.
`http://www.sigdial.org/workshops/conference18/proceedings/index.html`.

Semantic Inferencing, Concept Matching, Question Answering. [Chen *et al.*, 2010; Branavan *et al.*, 2012; Liang *et al.*, 2013; Tsai and Roth, 2016; Bakhshandeh and Allen, 2017; Richardson and Kuhn, 2017]

BioNLP, NLP for health sciences [Angelova and Boytcheva, 2011; Raghavan *et al.*, 2012; Mayfield *et al.*, 2013; Hassan *et al.*, 2015; Althoff *et al.*, 2016; Pérez-Rosas *et al.*, 2017]
Papers from workshops on biomedical NLP: `https://aclweb.org/anthology/W/W15/#3800`
`http://aclweb.org/anthology/sigbiomed.html#2017_0`

Other Innovative Applications NL and image / video retrieval: [Zarrieß and Schlangen, 2017]; papers from e.g. ACM SIGMM conferences `http://www.sigmm.org/view_conference`
NL interfaces for educational applications: [Song *et al.*, 2017]; papers from workshops on NLP for Educational Applications: `https://aclanthology.coli.uni-saarland.de/events/ws-2017#W17-50` `https://aclanthology.coli.uni-saarland.de/events/ws-2018#W18-05`
Educational Data Mining Conferences `http://educationaldatamining.org/`

References

[Agarwal *et al.*, 2013] Apoorv Agarwal, Anup Kotalwar, and Owen Rambow. Automatic extraction of social networks from literary text: A case study on Alice in Wonderland. In *Proceedings of the 6th International Joint Conference on Natural Language Processing (IJCNLP 2013)*, Nagoya, Japan, 2013.

[Althoff *et al.*, 2016] Tim Althoff, Kevin Clark, and Jure Leskovec. Large-scale analysis of counseling conversations: An application of natural language processing to mental health. *Transactions of the Association for Computational Linguistics*, 4:463–476, 2016.

[Angelova and Boytcheva, 2011] Galia Angelova and Svetla Boytcheva. Towards temporal segmentation of patient history in discharge letters. In *Proceedings of the Second Workshop on Biomedical Natural Language Processing*, pages 49–54, Hissar, Bulgaria, September 2011.

[Bakhshandeh and Allen, 2017] Omid Bakhshandeh and James Allen. Apples to apples: Learning semantics of common entities through a novel comprehension task. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 906–916, 2017.

[Bracewell, 2015] David Bracewell. Long nights, rainy days, and misspent youth: Automatically extracting and categorizing occasions associated with consumer products. In *Proceedings of the third International Workshop on Natural Language Processing for Social Media*, pages 29–38, Denver, Colorado, June 2015. Association for Computational Linguistics.

[Branavan *et al.*, 2012] S.R.K. Branavan, Nate Kushman, Tao Lei, and Regina Barzilay. Learning high-level planning from text. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 126–135, Jeju Island, Korea, July 2012. Association for Computational Linguistics.

[Chen and Di Eugenio, 2013] Lin Chen and Barbara Di Eugenio. Multimodality and Dialogue Act Classification in the RoboHelper Project. In *Proceedings of the SIGDIAL 2013 Conference*, pages 183–192, Metz, France, August 2013.

[Chen *et al.*, 2010] D.L. Chen, J. Kim, and R.J. Mooney. Training a multilingual sportscaster: Using perceptual context to learn language. *Journal of Artificial Intelligence Research*, 37(1):397–436, 2010.

[Di Eugenio *et al.*, 2013] Barbara Di Eugenio, Nick Green, and Rajen Subba. Detecting Life Events in Feeds from Twitter. In *ICSC 13, Seventh International Conference on Semantic Computing*, pages 274–277. IEEE, 2013.

[Fang *et al.*, 2014] Rui Fang, Malcolm Doering, and Joyce Y. Chai. Collaborative models for referring expression generation in situated dialogue. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 1544–1550, 2014.

[Ghosh *et al.*, 2017] Debanjan Ghosh, Alexander Richard Fabbri, and Smaranda Muresan. The role of conversation context for sarcasm detection in online interactions. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 186–196, 2017.

[Goldwasser and Zhang, 2016] Dan Goldwasser and Xiao Zhang. Understanding satirical articles using common-sense. *Transactions of the Association for Computational Linguistics*, 4:537–549, 2016.

[Hassan *et al.*, 2015] Mohsen Hassan, Olfa Makkaoui, Adrien Coulet, and Yannick Toussain. Extracting disease-symptom relationships by learning syntactic patterns from dependency graphs. In *Proceedings of BioNLP 15*, pages 71–80, Beijing, China, July 2015. Association for Computational Linguistics.

[Janarthanam *et al.*, 2013] Srinivasan Janarthanam, Oliver Lemon, Xingkun Liu, Phil Bartie, William Mackaness, and Tiphaine Dalmas. A multithreaded conversational interface for pedestrian navigation and question answering. In *Proceedings of the SIGDIAL 2013 Conference*, pages 151–153, Metz, France, August 2013. Association for Computational Linguistics.

[Johnson and Goldwasser, 2018] Kristen Johnson and Dan Goldwasser. Classification of moral foundations in microblog political discourse. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 720–730. Association for Computational Linguistics, 2018.

[Liang *et al.*, 2013] Percy Liang, Michael I. Jordan, and Dan Klein. Learning dependency-based compositional semantics. *Computational Linguistics*, 39(2), 2013.

[Mairesse and Walker, 2010] F. Mairesse and M.A. Walker. Towards Personality-based User Adaptation: Psychologically-informed Stylistic Language Generation. *User Modeling and User-Adapted Interaction*, 20(3):227–278, 2010.

[Manuvinakurike *et al.*, 2016] Ramesh Manuvinakurike, Maike Paetzel, Cheng Qu, David Schlangen, and David DeVault. Toward incremental dialogue act segmentation in fast-paced interactive dialogue systems. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 252–262, 2016.

[Mayfield *et al.*, 2013] Elijah Mayfield, David Adamson, and Carolyn Penstein Rosé. Recognizing rare social phenomena in conversation: Empowerment detection in support group chatrooms. In *ACL*, pages 104–113, 2013.

[Misra *et al.*, 2017] Amita Misra, Shereen Oraby, Shubhangi Tandon, Pranav Anand, and Marilyn Walker. Summarizing dialogic arguments from social media. In *Proc. SEMDIAL 2017 (SaarDial) Workshop on the Semantics and Pragmatics of Dialogue*, pages 126–136, 2017.

[Pavlick and Tetreault, 2016] Ellie Pavlick and Joel Tetreault. An empirical analysis of formality in online communication. *Transactions of the Association for Computational Linguistics*, 4:61–74, 2016.

[Pérez-Rosas *et al.*, 2017] Verónica Pérez-Rosas, Rada Mihalcea, Kenneth Resnicow, Satinder Singh, and Lawrence An. Understanding and predicting empathic behavior in counseling therapy. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1426–1435, 2017.

[Prabhakaran and Rambow, 2013] Vinodkumar Prabhakaran and Owen Rambow. Written dialog and social power: Manifestations of different types of power in dialog behavior. In *Proceedings of the 6th International Joint Conference on Natural Language Processing (IJCNLP 2013)*, Nagoya, Japan, 2013.

[Preoţiuc-Pietro *et al.*, 2017] Daniel Preoţiuc-Pietro, Ye Liu, Daniel Hopkins, and Lyle Ungar. Beyond binary labels: Political ideology prediction of twitter users. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 729–740, Vancouver, Canada, July 2017. Association for Computational Linguistics.

[Raghavan *et al.*, 2012] Preethi Raghavan, Eric Fosler-Lussier, and Albert Lai. Temporal Classification of Medical Events. In *BioNLP: Proceedings of the 2012 Workshop on Biomedical Natural Language Processing*, pages 29–37, Montréal, Canada, June 2012. Association for Computational Linguistics.

[Richardson and Kuhn, 2017] Kyle Richardson and Jonas Kuhn. Learning semantic correspondences in technical documentation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1612–1622, 2017.

[Shalyminov *et al.*, 2017] Igor Shalyminov, Arash Eshghi, and Oliver Lemon. Challenging neural dialogue models with natural data: Memory networks fail on incremental phenomena. In *Proc. SEMDIAL 2017 (SaarDial) Workshop on the Semantics and Pragmatics of Dialogue*, pages 98–106, 2017.

[Singh *et al.*, 2002] Satinder Singh, Diane Litman, Michael Kearns, and Marilyn Walker. Optimizing Dialogue Management with Reinforcement Learning: Experiments with the NJFun System. *Journal of Artificial Intelligence Research (JAIR)*, 15:105–133, 2002.

[Song *et al.*, 2017] Wei Song, Dong Wang, Ruiji Fu, Lizhen Liu, Ting Liu, and Guoping Hu. Discourse mode identification in essays. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 112–122, Vancouver, Canada, July 2017. Association for Computational Linguistics.

[Tsai and Roth, 2016] Chen-Tse Tsai and Dan Roth. Concept grounding to multiple knowledge bases via indirect supervision. *Transactions of the Association for Computational Linguistics*, 4:141–154, 2016.

[Zarrieß and Schlangen, 2017] Sina Zarrieß and David Schlangen. Obtaining referential word meanings from visual and distributional information: Experiments on object naming. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 243–254, Vancouver, Canada, July 2017. Association for Computational Linguistics.

[Zhang *et al.*, 2018] Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213. Association for Computational Linguistics, 2018.